arXiv:1608.02800v1 [cs.PF] 9 Aug 2016

LITMUS: An Open Extensible Framework for Benchmarking RDF Data Management Solutions

Harsh Thakkar[1], Mohnish Dubey[1], Gezim Sejdiu[1], Axel-Cyrille Ngonga Ngomo[2], Jeremy Debattista[1,3], Christoph Lange[1,3], Jens Lehmann[1,3], Sören Auer[1,3], Maria-Esther Vidal[1,3]

[1] University of Bonn;
{thakkar,dubey,sejdiu,debattis,langec,jens.lehmann,auer,vidal}@cs.uni-bonn.de
[2] University of Leipzig;
ngonga@informatik.uni-leipzig.de
[3] Fraunhofer IAIS;

Abstract. Developments in the context of Open, Big, and Linked Data have led to an enormous growth of structured data on the Web. To keep up with the pace of efficient consumption and management of the data at this rate, many data management solutions have been developed for specific tasks and applications. We present LITMUS, a framework for benchmarking data management solutions. LITMUS goes beyond classical storage benchmarking frameworks by allowing for analysing the performance of frameworks across query languages. In this position paper we present the conceptual architecture of LITMUS as well as the considerations that led to this architecture.

1 Introduction

Vast amounts of structured (following Linked Data principles) and un/semi-structured data is constantly being made available on the Web, often in an open manner[4], and within organisations. This rapid growth of data, available across organisations, has affected the data management layer of modern applications. Consequently, organisations are increasingly facing the need to find data management tools suited for the specific tasks at the core of their information management. Choosing the best data management tool is, however, challenging due to the limited comparability and compatibility of existing evaluation results and benchmarks. With regard to the limited domain expertise of the end user, the need for standardised frameworks to benchmark and analyse the existing diverse data management platforms is consequently of paramount importance.

Despite the growing interest and use in both research and the industry communities, currently the creators of benchmarks for Data Management Solutions (DMS) [1,3] do not offer a common suite for performing cross-domain benchmarks (i.e. one-to-one comparison of RDF, Graph, Wide-column, Relational, etc stores). In addition, there is no significant baseline to compare these cross-domain DMSs one against the other. Moreover, reproducing benchmarks is a

[4] With *open* we follow the Open Data Definition http://opendefinition.org

non-trivial problem owing to reasons such as non-standardised setup configurations, lack of publicly available resources (such as scripts, libraries, packages, etc.) and lack of transparent evaluation policies. Results in areas such as named entity recognition and linking [16] as well as question answering [14,15] have, however, shown that the provision of standardised interfaces and measures can contribute to the improvement of the performance of software solutions.

In this position paper we present the concept behind LITMUS, an open extensible approach for benchmarking a wide variety of DMS for storing RDF. LITMUS aims to provide support to organisations aspiring to use Linked Data management technologies in a wide spectrum of applications and magnitudes. LITMUS will provide a realistic performance evaluation platform covering a plethora of heterogeneous technologies (see Section 4) for storage and query benchmarking. To put the reader into the context of this work, and to highlight the objectives of LITMUS, we present the following **user scenario**:

"The WDAqua research project[5] aims towards building a data-driven question answering platform by using Web data, available in various formats, e.g., RDF, CSV, SQL, or XML. Harsh, a researcher within the project, is responsible for ensuring efficient data management (storage and retrieval) for this project. There is a large number of DMSs, each deliberately tailored to handling specific formats of data and queries, which need to be benchmarked to select the best solution for the project's needs. However, benchmarking of DMSs is non-trivial: it takes large amounts of human effort in designing, administering, evaluating, and analysing the diverse systems involved. Additionally, for the research project, a large set of factors, e.g., query typology, indexing speed, index size, query response time, and dataset size, need to be considered to ensure reproducibility and generality of the observed experimental results. *Harsh, wants to automate the whole benchmarking process, allowing easy integration, evaluation on custom stress loads, and fast analysis of the evaluation results.* He would also *expect the framework to be flexible to integrate new DMSs to the plethora of existing systems and benchmark them against a baseline.* Thus, Harsh's research question is: *Can a computational framework provide the required support for identifying the independent factors of his experiments, and for analysing and interpreting of the experimental results?*. The answer to this research question is *yes*, and the computational framework is LITMUS, an open extensible platform for benchmarking cross-domain DMSs. LITMUS will not only satisfy Harsh's need for automating the tedious benchmarking process, but will also offer: ***(1)*** an efficient way for replicating existing benchmarks (e.g., BSBM [3] or WAT-DIV [1]); ***(2)*** a wide set of performance evaluation measures/indicators tailored specifically for the DMS being evaluated; and ***(3)*** the comparison and visualisation of the performance of benchmarked DMSs on various intrinsic factors via custom charts, graphs and tabular data."

The remainder of this article is organised in the following sections: (2) *Related work* on benchmarking efforts, and their shortcomings, (3) *Objectives, Challenges and Outcomes*, which shed light on the focus of LITMUS, (4) *Framework*, describing the components the LITMUS components, and (5) *Conclusions*, summarizing the article.

[5] WDAqua ITN – http://wdaqua.eu

2 Related work

Benchmarking is widely used for evaluating data stores. Benchmarks exist for a variety of levels of abstraction from simple data models to graphs and triple stores, to entire enterprise information systems. We describe the current state of the art in benchmarking, in particular benchmarks for (a) relational databases, (b) graph databases, (c) RDF stores, (d) key-value stores, (e) wide-column stores, and (f) cross-domain benchmarking efforts. We identify shortcomings and limitations of existing systems, in order to determine the gaps that LITMUS needs to take into consideration. In addition to surveying existing work, we intend to focus mainly on the purpose and scope of the benchmarks.

In *Relational* DMSs, the benchmarks of the Transaction Processing Performance Council (**TPC**) [10] are well established. TPC uses discrete metrics for measuring the performance of the relational DMS. The online transaction processing benchmarks **TPC-C** and **TPC-E** use a transactions per minute metric. The analytics **TPC-H** and decision support **TPC-DS** benchmarks use the queries per hour and cost per performance metrics respectively.

For benchmarking *Graph* DMS, there are some existing works in their early stages (such as **HPC** Scalable Graph Analysis Benchmark [5], **Graph 500** [9], **XGDBench** [4]) dealing with graph suitability transformations and graph analysis. However they fail to define standards for graph modeling and query languages.

The substantial increase in the number of applications that use *RDF* data has encouraged the need for large scale benchmarking efforts on all aspects of the Linked Data life cycle, mostly focusing on query processing [11]. RDF DMS benchmarks make use of real (i.e., DBpedia or Wikidata) and synthetic (i.e., Berlin SPARQL Benchmark or WAT-DIV) datasets to evaluate DMS performance over custom stress-loads and setup environments [12].[6] DBpedia SPARQL Benchmark **(DBPSB)** [8] assesses RDF DMSs performance over DBpedia by creating a query workload derived from the DBpedia query logs.

The aim of the Lehigh University Benchmark (**LUBM**) [7] is to evaluate the performance of Semantic Web triple stores over a large synthetic dataset that complies to a university domain ontology. The Berlin SPARQL Benchmark (**BSBM** [3]) is another benchmark based on synthetic data, which addresses e-commerce use cases built around a set of products offered by different vendors. The Waterloo SPARQL Diversity TEST Suite (**WatDiv** [1]), provides data and query generators to enable benchmarking of RDF DMSs against a varying query structure (also complexity) to understand correlation of query typology with the variance in DMS performance. **SP2Bench** [13], one of the most commonly used synthetic data based benchmarks, uses the schema of the DBLP bibliographic dataset[7] to generate arbitrarily large datasets.

[6] https://www.w3.org/wiki/RdfStoreBenchmarking
[7] http://dblp.uni-trier.de/db/

There are only a few efforts that benchmark *cross-domain* DMS. **Pandora**[8], one such effort, uses the Berlin SPARQL Benchmark data to benchmark RDF stores against relational stores (Jena-TDB, Monetdb, GH-RDF-3X, PostgreSQL, 4Store). **Graphium** [6] is a similar study benchmarking RDF stores against Graph stores (Neo4J, Sparksee/DEX, HypergraphDB, RDF-3X) on graph datasets including a 10M triple graph data generated using the Berlin SPARQL Benchmark data generator. More recently, the **LDBC** [2] focuses on combining industry-strength benchmarks for graph and RDF data management systems. The LDBC introduces a new bottleneck methodology for developing benchmark workloads, which tries to combine user input with feedback from system experts.

Research has so far focused on benchmarking domain specific DMSs, despite the need for integrating cross-domain DMSs and automating the benchmarking process. LITMUS aims at addressing these shortcomings and serving as an open, extensible platform to allow easy integration, benchmarking and performance analysis of diverse data management solutions. To the best of our knowledge, no such open, extensible and reusable framework exists, which allows to explore and analyse a wide range of different DMSs.

3 Objectives, Challenges and Outcomes

3.1 Focus of the LITMUS framework

The LITMUS framework aims at bridging the gaps in adopting, deploying and scaling the consumption of Linked Data. LITMUS focuses on simplifying the use, assessment and analysis of the performance of a wide spectrum of cross-domain DMSs. In particular, the LITMUS project will:

- **F1** enable a common ground for benchmarking and comparing a plethora of cross-domain DMSs, and replicating existing third-party benchmarks;
- **F2** create *(i)* interoperable machine-readable evaluation reports and *(ii)* scientific studies on the correlation of a variety of factors (such as query typology, data structures used for indexing, etc.) with respect to the performance of DMSs;
- **F3** recommend particular DMSs and benchmarks based on a set of user predefined requirements.

3.2 Challenges to be addressed

To develop such an open extensible benchmarking platform, three key challenges have to be addressed:

- **C1** **Data conversion**: This challenge demands a generic data conversion mechanism allowing users to convert the RDF data to a format interpretable by the corresponding DMS. The focus is to represent RDF data in multiple formats, keeping the end user as secluded as possible from the framework's technical details.

[8] http://pandora.ldc.usb.ve/

– **C2 Query Conversion:** Cross-domain benchmarking of DMSs demands that queries are represented in all languages and formats supported by the respective tools. Query languages differ in their structure and expressivity. For instance, complex path queries (in SPARQL, in particular Kleene stars) cannot be expressed in an equivalent SQL query. There is a need to develop an intermediate mechanism to convert or express the logic of one query (e.g. form SPARQL) to the other respective language (e.g. to CYPHER, SQL, CQL). This requires an exhaustive study of the query languages' specifications. The main challenge is to identify the correct mappings between different languages, maintaining the correctness and meaning of the original query.
– **C3 Performance indicators:** The performance of a DMS can be assessed with regard to a variety of indicators. Dealing with the diverse characteristics of the DMSs, it is necessary to explore complex performance indicators in contrast to traditional ones, namely precision, recall, index size, storage size, number of triples, and query response time.

3.3 Outcomes of LITMUS

The artifacts resulting from the LITMUS project will be *(A1)* scientific studies and *(A2)* frameworks/software.

A1*Scientific studies:*

– An in-depth analysis of the query language expressivity and supported features striving to address the language barrier **(C2)** (ref. section 3.2). This study will provide us with deep insights about the functionality of various query languages, their strengths and limitations.
– An exhaustive exploratory study on the selection of performance measures for evaluating cross-domain DMSs, addressing challenge**(C3)**(section 3.2)

A2 *Framework/Software (i.e. algorithms, tools, etc):*

– Automatic conversion of RDF data to multiple data formats (such as CSV, JSON, SQL, etc.), providing compatible data as input to the cross-domain DMSs.
– Novel mechanisms for the automatic conversion of SPARQL to format-specific query languages, enabling compatible query input for cross-domain DMSs.
– An open, extensible benchmarking platform, LITMUS, for cross-domain DMS performance evaluation and easy replication of existing benchmarks.

3.4 Target audience

Technology Vendors: This addresses developers of commercial, industrial DMSs (including system and data analysts, system developers, system architects) who are thriving towards developing more and more advanced DMSs for efficient consumption of Big Data.
Technology Consumers: Staff of private and commercial organisations and

other users seeking recommendations for the best solution for their needs can simply compare a wide range of DMS against a list of desired parameters.
Technology Researchers: The researchers who can benefit from our lessons-learned , and researchers whom LITMUS enables to contribute further results to the community. Target communities include Semantic Web, Databases, Information Retrieval, Big Data and others.

4 The Litmus Framework

4.1 Architecture Overview

The architecture of the LITMUS framework will comprise four major facets: Data Facet (F1), Query Facet (F2), System Facet (F3) and Benchmarking Core (F4) (Figure 1) In the following, we explain the role of each facet.

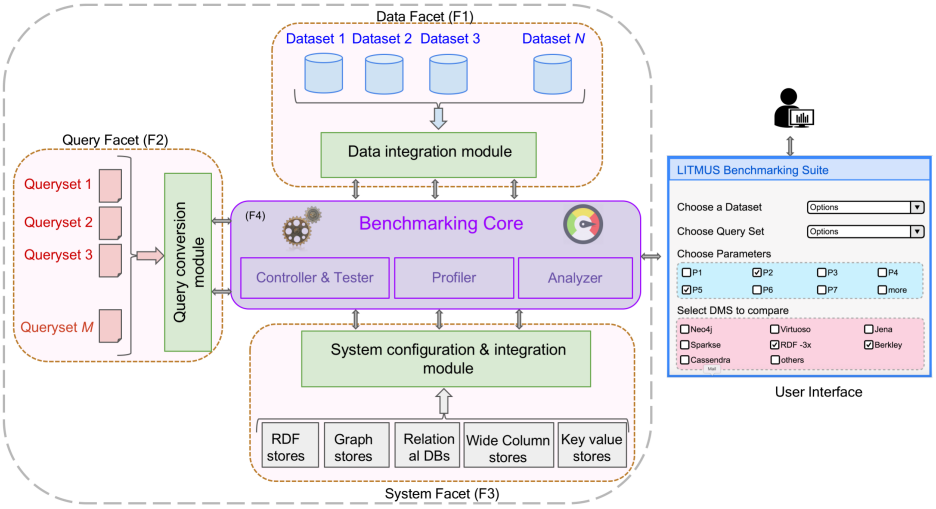

Fig. 1. Overview of the LITMUS framework architecture.

Data Facet F1: The Data Facet would deal with Dataset(s) and the Data Integration Module. Datasets chosen for benchmarking may be real datasets such as DBpedia[9] Wikidata,[10]; synthetic datasets such as the Berlin SPARQL Benchmarking (BSBM) [3], Waterloo SPARQL Diversity Test Suite (WatDiv) [1], or hybrid datasets comprising both real and synthetic data. The *Data Integration Module* is responsible for (a) making data available to the system in the requested formats (such as N-Triples, CSV, SQL or JSON) by carrying out appropriate

[9] http://wiki.dbpedia.org/
[10] https://www.wikidata.org

data conversion and mapping tasks (cf. Challenge **C1**), and (b) loading the desired format of data to the respective DMSs selected for the benchmark.

Query Facet F2: The Query Facet would deal with Queryset(s), and the Query Conversion Module. The *Queryset* refers to the set of query input files. The *Query Conversion Module* will be one of the key components addressing the language barrier (Challenge **C2**). It is responsible for converting the input SPARQL queries to the respective DMSs' query languages (such SQL, CYPHER or CQL). The conversion will be performed by developing an intermediate language/logic representation of the input query. The aim of this module is to allow efficient conversion of a wide variety of SPARQL queries (such as path, star-shaped and snowflake queries) to other query languages, ultimately breaking the language barrier.

System Facet F3: The System Facet also consists of two key modules, (i) DMSs and (ii) DMS Configuration and Integration module. The *DMSs* module consists of the DMS selected for the benchmark. The *DMS Configuration and Integration* module is responsible for (i) providing easy integration, via wrapper(s) or as a plugin, of the DMS, and (ii) monitor and configure the the integrated DMS for the benchmark. On top of this, this module will make use of *Docker*[11] *containers* to ensure a fair allocation of resources and to provide the necessary isolation required for conducting realistic benchmarks.

Benchmarking Core F4: The Benchmarking Core is the heart of the LITMUS framework, consisting of three modules: (i) Controller and Tester, (ii) Profiler, and (iii) Analyser. The *Controller and Tester* is responsible for executing the respective scripts for loading data and fetching the queries to their corresponding DMSs, creating and validating the specified system configurations, and finally executing the benchmark on the selected setting. The *Profiler* is responsible for: (a) generating and loading various profiles (stress loads, query variations, etc.) for conducting the benchmark tests and (b) storing the benchmark results profile-wise. The *Analyser* is responsible for collecting the benchmark results from the *Profiler* and generates performance evaluation reports. It also performs a correlation analysis between the parameters specified by the user. The final results (reports) will then presented to the end user in a suitable visualisation.

5 Conclusions

LITMUS addresses the gaps of the cross-benchmarking platform for different query languages and corresponding data management solutions. The literature review confirms the absence of such a cross-benchmarking platform. We have mentioned the upcoming challenges, which the proposed system will have to address. The proposed architecture of LITMUS would provide solutions to these challenges.

Acknowledgements: Parts of this work have been supported by the EU Horizon 2020 Framework Programme under grant agreements no. 642795 (WDAqua ITN), 644564 (Big Data Europe) and 688227 (HOBBIT).

[11] Docker – https://www.docker.com/

References

1. G. Aluç, O. Hartig, M. T. Özsu, and K. Daudjee. Diversified stress testing of RDF data management systems. In *International Semantic Web Conference*. Springer, 2014.
2. R. Angles, P. A. Boncz, J. Larriba-Pey, et al. The linked data benchmark council: a graph and RDF industry benchmarking effort. *SIGMOD Record*, 2014.
3. C. Bizer and A. Schultz. The berlin SPARQL benchmark. *Int. J. Semantic Web Inf. Syst.*, 5, 2009.
4. M. Dayarathna and T. Suzumura. XGDBench: A benchmarking platform for graph stores in exascale clouds. In *CloudCom*. IEEE Computer Society, 2012.
5. D. Dominguez-Sal, P. Urbón-Bayes, A. Giménez-Vañó, et al. Survey of graph database performance on the HPC scalable graph analysis benchmark. In *Proceedings of the 2010 International Conference on Web-age Information Management*, WAIM'10. Springer-Verlag, 2010.
6. A. Flores, G. Palma, M.-E. Vidal, et al. GRAPHIUM: visualizing performance of graph and rdf engines on linked data. In *Proceedings of the 2013th International Conference on Posters & Demonstrations Track-Volume 1035*. CEUR-WS. org, 2013.
7. Y. Guo, Z. Pan, and J. Heflin. LUBM: A benchmark for owl knowledge base systems. *Web Semant.*, 3, Oct. 2005.
8. M. Morsey, J. Lehmann, S. Auer, and A.-C. Ngonga Ngomo. *DBpedia SPARQL Benchmark – Performance Assessment with Real Queries on Real Data*. Springer Berlin Heidelberg, 2011.
9. R. C. Murphy, K. B. Wheeler, B. W. Barrett, and J. A. Ang. Introducing the GRAPH 500. *Cray User's Group (CUG)*, 2010.
10. R. Nambiar, N. Wakou, F. Carman, and M. Majdalany. *Transaction Processing Performance Council (TPC): State of the Council 2010*. Springer Berlin Heidelberg, 2011.
11. A.-C. N. Ngomo and M. Röder. HOBBIT: Holistic benchmarking for big linked data. *ERCIM News*, 2016.
12. M. Saleem, Y. Khan, A. Hasnain, I. Ermilov, and A. N. Ngomo. A fine-grained evaluation of SPARQL endpoint federation systems. *Semantic Web*, 7, 2015.
13. M. Schmidt, T. Hornung, M. Meier, C. Pinkel, and G. Lausen. SP2Bench: A SPARQL performance benchmark. In *Semantic Web Information Management*. Springer, 2009.
14. G. Tsatsaronis, G. Balikas, P. Malakasiotis, et al. An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, 16, 2015.
15. C. Unger, C. Forascu, V. Lopez, et al. Question answering over linked data (QALD-5). In *Working Notes of CLEF 2015, Toulouse, France, September 8-11*, 2015.
16. R. Usbeck, M. Röder, A. N. Ngomo, et al. GERBIL: general entity annotator benchmarking framework. In *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22*, 2015.